WHAT IS CLAIMED IS:

1. An audio processing method, comprising:

sequentially rendering audio summaries and transition audio segments with at least one transition audio segment rendered between each pair of sequential audio summaries, wherein each audio summary comprises digital content summarizing at least a portion of a respective associated audio piece.

2. The method of claim 1, wherein identical transition audio segments are rendered between pairs of sequential audio summaries.

3. The method of claim 2, wherein each identical transition audio segment corresponds to a Gabor function in a time domain representation.

4. The method of claim 3, wherein each Gabor function has a center frequency substantially corresponding to a center pitch of an adjacent audio summary.

5. The method of claim 1, wherein the audio summaries and the interleaved transition audio segments are rendered consecutively.

6. The method of claim 1, wherein each audio summary is a representative segment of a respective associated audio piece.

7. The method of claim 1, further comprising classifying audio pieces into categories in response to user input received during rendering of the associated audio summaries.

8. The method of claim 7, further comprising building a playlist based on categories assigned to a set of audio pieces.

9. The method of claim 1, wherein at least one audio summary is linked to an associated audio piece.

10. The method of claim 9, further comprising rendering an audio piece linked to an associated audio summary in response to user input received during rendering of the associated audio summary.

11. The method of claim 1, further comprising rendering a given audio piece beginning at a location in the given audio piece linked to an audio summary associated with the given audio piece.

12. The method of claim 11, further comprising rendering a second audio piece at a location in the second audio piece linked to a successive audio summary associated with the second audio piece.

13. The method of claim 1, further comprising ordering audio summaries in a sequence based on similarity to a given audio summary.

14. The method of claim 13, wherein audio summaries are rendered in accordance with the ordered sequence.

15. The method of claim 1, wherein each audio piece is associated with multiple audio summaries and a single audio summary is rendered automatically for each audio piece, and further comprising rendering an audio summary for a given audio piece in response to user input received during rendering of a preceding audio summary associated with the given audio piece.

16. The method of claim 1, further comprising normalizing audio summaries to a common loudness level.

17. An audio processing system, comprising:

a rendering engine operable to sequentially render audio summaries and transition audio segments with at least one transition audio segment rendered between each pair of sequential audio summaries.

18. A method of generating an annotated audio file, comprising:

annotating an original audio file by embedding therein information enabling rendering of at least one audio summary contained in the annotated audio file and comprising digital content summarizing at least a portion of the original audio file.

19. The method of claim 18, wherein the rendering enabling information is embedded in a header of the audio file.

20. The method of claim 19, wherein rendering enabling information includes an audio summary embedded in the audio file header.

21. The method of claim 19, wherein rendering enabling information embedded in the audio file header includes one or more pointers to one or more respective locations in the original audio file.

22. The method of claim 18, wherein rendering enabling information is embedded at different locations in the annotated audio file separated by audio content of the original audio file.

23. The method of claim 22, wherein rendering enabling information includes audio summaries embedded at different respective locations in the annotated audio file separated by audio content of the original audio file.

24. The method of claim 22, wherein rendering enabling information includes pointers to locations in the original audio file, the pointers being embedded at different respective locations in the annotated audio file separated by audio content of the original audio file.

25. The method of claim 18, wherein rendering enabling information includes hierarchical information enabling rendering of audio summaries at different levels of detail.

26. The method of claim 18, wherein at least one audio summary corresponds to a representative sample of the original audio file.

27. The method of claim 18, wherein at least one audio summary corresponds to audio content not contained in the original audio file.

28. The method of claim 18, wherein at least one audio summary corresponds to one or more images representative of original audio file content.

29. The method of claim 18, wherein at least one audio summary corresponds to digital textual content.

30. A software program for generating an annotated audio file, the software program residing on a medium readable by an electronic device and comprising instructions for causing an electronic device to:

annotate an original audio file by embedding therein information enabling rendering of at least one audio summary contained in the annotated audio file and comprising digital content summarizing at least a portion of the original audio file.

31. A method of generating an annotated audio file, comprising:

annotating an original audio file by providing at least one browsable link between the original audio file and at least one audio summary comprising digital content summarizing at least a portion of the original audio file, and storing the original audio file, the at least one browsable link, and the at least one audio summary on a common portable storage medium.

32. A portable medium readable by an electronic device and tangibly storing an original audio file, at least one audio summary comprising digital content summarizing at least a portion of an original audio file, and at least one browsable link between the original audio file and the at least one audio summary.

33. A system for rendering an annotated video file, comprising:

a rendering engine operable to identify information embedded in the annotated audio file and enabling rendering of at least one audio summary contained in the annotated audio file and comprising digital content summarizing at least a portion of the original audio file, and to operable to render the at least one audio summary.

34. An audio processing method, comprising:

dividing an audio piece into audio segments;

extracting acoustical features from each audio segment;

grouping audio segments into clusters based on the extracted features;

identifying a representative audio segment in each cluster; and

selecting a representative audio segment as an audio summary of the audio piece.

35. The method of claim 34, wherein each audio segment has a substantially equal rendering time.

36. The method of claim 34, further comprising computing a feature vector centroid for each cluster, wherein each representative audio segment is closer to the feature vector centroid computed for the corresponding cluster than all other audio segments in the corresponding cluster.

37. The method of claim 34, further comprising ranking clusters based on respective numbers of audio segments in the clusters.

38. The method of claim 37, wherein a representative audio segment of a highest ranked cluster is selected as the audio summary.